



(12)

## EUROPEAN PATENT APPLICATION

(43) Date of publication:  
27.11.2002 Bulletin 2002/48

(51) Int Cl.7: H04L 12/56

(21) Application number: 02291078.0

(22) Date of filing: 29.04.2002

(84) Designated Contracting States:  
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU  
MC NL PT SE TR  
Designated Extension States:  
AL LT LV MK RO SI

(30) Priority: 24.05.2001 US 864971

(71) Applicant: ALCATEL  
75008 Paris (FR)

(72) Inventors:  
• Basso, Claude  
Raleigh NC 27613 (US)  
• Arts, Francis Luc Mathilda  
2370 Arendonk (BE)

- Debuyscher, Pierre Leon  
9840 De Pinte (BE)
- Duroyon, Olivier Didier  
San Diego, CA 92120 (US)
- Povse, Max Robert  
Apex, North Carolina 27502 (US)
- Vaidhyanathan, Natarajan  
Greensboro, NC 27407 (US)
- Verilli, Colin  
Apex, North Carolina 27509 (US)


(74) Representative:  
Narmon, Gisèle Marie Thérèse et al  
Alcatel Bell N.V.,  
Francis Wellesplein 1  
2018 Antwerpen (BE)

(54) **System and method for enhancing the availability of routing systems through equal cost multipath**

(57) In a networking environment including one or more network processing (NP) devices and implementing a routing protocol for routing data packets from a source NP devices to destination NP devices via a switch fabric, with each network processing device supporting a number of interface ports, a system and method for enabling a routing system to recover more quickly that the routing protocol so as to significantly reduce the occurrence of lost data packets to a failed target interface/blade. The routing system is enabled to track the operational status of each network processor device and operational status of destination ports supported by

each network processor device in the system, and maintains the operational status as a data structure at each network processing device. Prior to routing packets, an expedient logical determination is made as to the operational status of a target network processing device and target interface port of a current packet to be routed as represented in the data structure maintained at the source NP device. If the target blade/interface is not operational, an alternative route may be provided by ECMP. In this manner, correct routing of packets is ensured with reduced occurrence of lost data packets due to failed target NP devices/ports.

FIG. 2A



the GPP, and, the programmable hardware-assist processors, picoocode in each of the network processors. These two network processor system 10 comprises two major software components: 1) the control point code base running on the GPP, and, the Network Processor (NP) devices so that they may handle the forwarding of data packets and frames. It should be understood however, that the GPP may be embedded in a network processor device itself. The generic router of the Network Processor (NP) devices so that they may handle the forwarding of data packets and frames, it is logical association with all of the Network Processor (GPP) functions as a control point (CP) 25 for the system and has a physical or separate General Purpose Processor (GPP) 12 in the system for enabling the customization and configuration of the Network Processor system 10. As shown in Figure 1, multiple Network Processors (NP) 12 are shown connected using a switch fabric 15, with each of the network processors supporting a large number of external LAN or WAN interface ports 20. A separate processor system 10. As shown in Figure 1, which illustrates a logical model of a generic Network Processor system 10. As shown in Figure 1, multiple Network Processors (NP) 12 are shown connected using a switch fabric 15, with each of the network processors supporting a large number of external LAN or WAN interface ports 20.

[0008] For exemplary purposes, reference is made to Figure 1, which illustrates a logical model of a generic Network Processor system 10.

55

50

45

40

35

30

25

20

15

10

5

## Discussion of the Prior Art

[0001] This invention relates generally to network processor-based devices, and more specifically to an improved equal cost multipath routing and recovery mechanism that enables the routing system to recover more quickly than the routing protocol.

[0002] In today's networked world, bandwidth is a critical resource. Increasing network traffic, driven by the Internet

## Background of the Invention

### Description

## Background of the Invention

software components are responsible for initializing the system, maintaining the forwarding paths, and managing the system. From a software view, the system is distributed. The GPP and each picoprocessor run in parallel, with the CP communicating with each picoprocessor using a predefined application program interface (API) 30 and control protocol.

[0009] The CP code base provides support for the Layer 2 and Layer 3 topology protocols and Layer 4 and Layer 5 network applications and systems management. Examples are protocol support for VLAN, IP, and Multiprotocol Label Switching standard (MPLS), and the supporting address- and route-learning algorithms to maintain topology information.

[0010] With particular reference to Figure 1, and accompanying description found in commonly-owned, co-pending U.S. Patent Application Serial No. 09/384,691 filed August 27, 1999 and entitled "NETWORK PROCESSOR PROCESSING COMPLEX AND METHODS", the whole contents and disclosure of which is incorporated by reference as if fully set forth herein, the general flow of a packet or frame received at the NP device is as follows: frames received from an network connection, e.g., Ethernet MAC, are placed in internal data store buffers by an upside "enqueue" device (EDS-UP) where they are identified as either normal data frames or system control frames (Guided Frames).

In the context of the invention, frames identified as normal data frames are enqueued to an Embedded Processor Complex (EPC) which comprises a plurality of picoprocessors, e.g., protocol processors. These picoprocessors execute logic (picocode) capable of looking at the received frame header and deciding what to do with the frame (forwardly, modify, filter, etc.). The EPC has access to several lookup tables, and classification hardware assists to allow the picoprocessors to keep up with the high-bandwidth requirements of the Network Processor. A classification hardware assist device in particular, is provided for classifying frames of well known frame formats. The Embedded Processing

Complex (EPC) particularly provides and controls the programmability of the NP device and includes, among other components (such as memory, dispatcher, interfaces), N processing units, referred to as GxH, which concurrently execute picocode that is stored in a common instruction memory. It is understood, however, that the architecture and structure is completely scalable towards more GxHs with the only limitation being the amount of silicon area provided in the chip. In operation, classification results from the classification hardware assist device are passed to the GxH,

25 during frame dispatch. Each GxH preferably includes a Processing Unit core (CLP) which comprises, e.g., a 3-stage pipeline, general purpose registers and an ALU. Several GxHs in particular, are defined as General Data Handlers (GDH) each of which comprise a full CLP with the five coprocessors and are primarily used for forwarding frames. One GxH coprocessor, in particular, a Tree Search Engine Coprocessor (TSE) functions to access all tables, counters, and other data in a control memory that are needed by the picocode in performing tree searches used in forwarding data 30 packets, thus freeing a protocol processor to continue execution. The TSE is particularly implemented for storing and retrieving information in various processing contexts, e.g., determining frame routing rules, lookup of frame forwarding information and, in some cases, frame alteration information.

[0011] Traditional frame routing capability provided in network processor devices typically utilize a network routing table having entries which provide a single next hop for each table entry. Commonly-owned, co-pending United States 35 Patent Application Serial No. 09/546,702 entitled METHOD FOR PROVIDING EQUAL COST MULTIPATH FORWARDING IN A NETWORK PROCESSOR, the whole content and disclosure of which is set forth herein, describes a system and method for providing the ability for a network processor to select from multiple next hop options for a single forwarding entry.

[0012] Figure 2(a) depicts an example network processor frame routing scenario 40 and Figure 2(b) illustrates an 40 example Equal Cost Multipath Forwarding (ECMP) table 50 that may be used to provide a lookup of a nextHop address for forwarding packets as described in commonly-owned, co-pending United States Patent Application Serial No. 09/546,702. Preferably, such a table is employed in a Network Processor (NP) device having packet routing functions such as described in commonly-owned, co-pending U.S. Patent Application 09/384,691.

[0013] Thus, the example ECMP forwarding table 50 illustrated in Figure 2(b), is particularly implemented in a frame 45 forwarding context for network processor operations. In the example ECMP forwarding table 50, there is provided subnet destination address fields 52, with each forwarding entry including multiple next hop routing information comprising multiple next hop address fields, e.g., fields 60a - 60c. Additionally provided in the ECMP routing table is cumulative probability data for each corresponding next hop such as depicted in action data field 70. Particularly, in the exemplary illustration of the ECMP packet forwarding table 50 of Figure 2(b), there is included three (3) next hop fields 50 to addresses 9.1.1.1, 8.1.1.1, 6.1.1.1 associated with a destination subnet address 7.\*.\*. An action data field 70 includes threshold values used to weight the probability of each next hop and is used to determine which next hop will be chosen. In the action field 72, shown in Figure 2(b), these values as being stored as cumulative percentages with the first cumulative percentage (30%) corresponding to next hop 0, the second cumulative percentage value (80%) corresponding to next hop 1, etc. This means that, the likelihood of routing a packet through next hop 0 is 30% (i.e., approximately 30% of traffic for the specified table entry should be routed to next hop 0), and, the likelihood of routing a packet through next hop 1 is 50% (i.e., approximately 50% of traffic for the specified table entry should be routed to next hop 1). This technique may be extended to offer as many next hops as desired or feasible.

[0014] Currently, in such network processing systems, if a destination NP device (hereinafter referred to as *Target*-

[Q02] A first method of maintaining operational status at the blade/NP level involves implementation of a data structure (hereinafter referred to as *status*) that is maintained by each NP device. This *status* data structure includes information representing the operational status of all the network processors (blades/ports) in the routing system and, for example, may comprises a bit vector of sixty-four (64) bits long (in an example system employing 64 NP devices).

## Detailed Description of the Preferred Embodiment

Figure 3 illustrates the determination of a failed link to Target Blade associated with EMCMP next hop destination NPI for the example network processing scenario of Figure 2(a), and the resulting decision to re-route the frame to an operation destination NPI according to the example EMCMP table.

Figure 2(b) illustrates an example Fowarding table for use in a network processor, router or packet switching device according to the example networking processing scenario of Figure 2(a).

Figure 2(a) depicts an example networking scenario 40 including network processors (routers) employing a packet routing table such as an EMC<sup>TM</sup> forwarding table.

Figure 1 illustrates a logical model of a generic Network Processor system 10.

[0019] Further features, aspects and advantages of the apparatus and methods of the present invention will become better understood with regard to the following description, appended claims, and accompanying drawings where:

### **Brief Description of the Drawings**

[Unit 8] Prior to routing packets, an expedient logical determination is made as to the operational status of a target network processing device and target interface port of a current packet to be routed as represented in the data structure maintained at the source NP device. In this manner, correct routing of packets is ensured with reduced occurrence of lost data packets due to failed target NP devices/ports.

structure at each network processing device.

including one or more network processing (NP) devices and implementing a routing protocol for routing data packets from a source to destination NP devices via a switch fabric, with each network processing device supporting a number of interface ports, a system and method for enabling a routing protocol for recovering more quickly than the routing protocol so as to significantly reduce the occurrence of lost data packets to a failed target interface/blade. The routing system is enabled to track the operational status of each network processor device and operational status of destination ports supported by each the system and maintains the operational status as a data structure in the system. The routing system is enabled to track the operational status of each network processor device and operational status of destination ports supported by each the system and maintains the operational status as a data structure in the system.

0017 In accordance with the preferred embodiment of the invention there is provided for a networking environment

It is another object of the present invention to provide in a network processor system, a method of maintaining the operational status of all the network processors (blades) in the routing system so that packet forwarding issues resulting from a failure of one or more blades can be resolved without loss of data packets sent in the system with minimal overhead.

[016] Accordingly, it is an object of the present invention to provide a network processor with a system that would enable a routing system to recover more quickly than the routing protocol so as to significantly reduce the oc-

### Summary of the Invention

**[0015]** Consequently, it would be highly desirable to provide a methodology that would enable a routing system to recover more quickly that the routing protocol so as to significantly reduce the occurrence of lost data packets to a failed target interface/blade with minimal performance penalty.

the failed link could be relatively long, and during this period all the data packets routed through the failed interface/blade. The time for this routing protocol to detect and download a new forwarding entry that avoids the failed interface/blade. and calculate the shortest path from an IP Source Address (SA) to IP Destination Address (DA), detects the failed link and calculate the shortest routes to understand the internal network architecture, i.e., within an autonomous network, protocol which enables routers to understand the routing protocol, e.g., the Open Shortest Path First (OSPF) forward frames through the failed interface/blade until the system will continue to attempt to forward frames through the other Network Processors (NPs) in the system set forth in the ECMP forwarding table. However, it is often the case that the correct destination is not found in the ECMP forwarding table. Therefore, the packet or frame cannot be routed to the target blade and capable of handling the frame type fails, i.e., the packet or frame associated with the target blade and capable of handling the blade or blade) or interface (such as a port or target-Port) associated with the target blade and capable of handling the failed link.

If the  $i$ th bit is set, for instance, then the  $i$ th NP/blade is indicated as operational. In operation, after choosing the next hop according to ECMP rules, the layer-3 forwarding picocode will check the operational status of the NP/blade through which the chosen next hop is reachable. If that NP is not operational, then a different equal-cost next hop (the next hop with the smallest index) that is reachable through an operational NP/blade will be chosen. Figure 3 illustrates the determination of a failed link to Target Blade associated with ECMP next hop destination NP1, and the resulting decision to re-route the frame to an operation destination NP2 according to the ECMP table. That is, in each NP, the operational status of the TB for each packet routed is checked. If the destination TB is down, then a different Next Hop is chosen as suggested by the ECMP table. It should be understood that the particular user application will detect failures and update the *opStatus* data structure accordingly.

[0021] This first solution essentially maintains the operational status at the TB (blade)/NP level. In order to extend this solution to an interface/port (TB/TP) level, there needs to be maintained a datastructure that is 64x16 bits long, assuming each blade in the example system maintains sixteen (16) ports, for instance. Since the *opStatus* datastructure is consulted in the main forwarding path, it must be stored in a fast, expensive memory.

[0022] Another solution relies on the assumption that the interface/blade failures are rare and it is unlikely that more than one blade will fail at the same time. The advantage of tracking a single failure is the reduction of the size of the *opStatus* data structure. The current solution only requires 48 bits in expensive high-speed memory where as the previous solution required 64 x 16 bits in such a memory. Thus, the following data structure may be maintained in each NP device in the routing system.

20

```

{
    Uint 16 failedBlade; /* Use the value of 0xffff if all blades are operational */
    Uint 16 failedPortMask;
    Uint 16 failedPortValue;
}
```

30

[0023] According to this embodiment, the following algorithm is invoked to check whether a given TB, TP is operational:

35

```

Boolean is Operational (TB, TP) {
    If (failedBlade == 0xffff)
        /* all blades are operational */
    return TRUE;
    If ((TB == failedBlade) && (TP & failedPortMask == failed PortValue))
```

40

```

    /* where && is the logical AND operator */
    /* where & is a bitwise AND operator*/
    Return FALSE;
```

45

Else

50

Return TRUE;

}

55

[0024] According to this algorithm, if all blades are operational, the routing of packets throughout the system will continue and no ECMP re-routing is necessary. However, only if both a Target Blade is a failed blade AND the result of the bitwise operation between the Target Port and failedPort Mask is equal to the failedPortValue, then a FALSE is returned and the ECMP table invoked for re-routing subsequent packets to another TB or TP. If a TRUE is returned,

[0029] According to this algorithm, if all blades are operational, then both failedPortMask and failedPortValue are set to 0xffffffff and the values of the other fields are ignored. This is a simple test that may be performed in one machine cycle.

55

```

    {
        Return FALSE;
    }

    If (TP & failedPortMask != failedPortValue) return TRUE;
    If (TB > endFailedBlade) return TRUE;
    If (TB < beginFailedBlade) return TRUE;
    return TRUE;
}

/* 1-cycle, 1 picocode instruction can perform this test */
/* all blades are operational */
If ((failedPortMask == 0xffffffff) && (failedPortValue == 0xffffffff))
    Boolean isOperational (TB, TP) {

```

35

[0028] According to this embodiment, the following algorithm is invoked to check whether a given TB, TP is operational. This convention is founded on the assumption that no port is numbered 0xffffffff. Considered operational. If failedPortMask and failedPortValue are both 0xffffffff, then all blades will be considered operational. This convention is founded on the assumption that no port is numbered 0xffffffff.

30

```

    {
        Unit failedPortValue;
        Unit 8 failedPortMasks;
        Unit 8 endFailedBlade; /* end value of range of failed blades */
        Unit 8 beginFailedBlade; /* unsigned integer representing begin value range of failed blades
        */
    }
}

Unit 8 beginFailedBlade; /* unsigned integer representing begin value range of failed blades

```

15

[0026] In the preferred embodiment, a range is used to represent the failed blades and a mask on the port number to represent the set of failed ports. This solution only requires 32 bits of high-speed memory. The following structure will be maintained in all of the NPs in the preferred embodiment:

10

failedBlade will contain the value of bladeNum, and failedPortMask will contain the value of 0x00003 and failedPortValue will contain the value of 0x0002. In the ports in DMU D have last 2 bits set to 1, if DMU C were to fail in blade numbered bladeNum, 2 bits set to 10, and the ports in DMU B have last 2 bits set to 01, the ports in DMU C have last 2 bits set (least significant) 2 bits set to 00, the ports in DMU A have last 2 bits set to 01, the ports in DMU A have the value of portNum. Assuming a blade having four data move units (DMUs) of four ports each, the failedBlade will contain bladeNum and failedPortMask will contain the value of 0xffffffff and the failedPortValue will contain the value of 0. If the port number failed bladeNum is not operational, then failedBlade will include bladeNum and failedPortMask will contain the value of 0 and failedPortValue will have failed) then failedBlade will be ignored. If the blade number, e.g., bladeNum, is not operational (i.e., all the ports in that blade failedPortValue will be ignored. If the blade number failed bladeNum is not operational, then failedBlade will contain the value of 0xffffffff and failedPortValue will be handled by this solution. As an example, if all the interfaces in all the blades are operational then failedBlade will contain the value of 0xffffffff and failedPortMask and blade levels. However, multiple blade failures cannot be handled by this solution. A data move unit (DMU) and failedPortMask is not equal to the failedPortValue, then the packet will still be routed to the destination TB/TR.

5

[0025] It should be understood that this solution may handle individual failures at port, data move unit (DMU) and

blade levels. However, multiple blade failures cannot be handled by this solution. As an example, if all the interfaces

If the blade numbered *bladeNum* is not operational (i.e., all the ports in that blade have failed) then, according to this algorithm,

beginFailedBlade and endFailedBlade are set as *bladeNum*,  
 failedPortMask is set as 0, and  
 failedPortValue is set as 0.

[0030] However, if the blades numbered, for example 8, 9, and 10 are not operational then set

beginFailedBlade as 8  
 endFailedBlade as 10  
 failedPortMask as 0 and  
 failedPortValue as 0

[0031] If the port numbered *portNum* in the blade numbered *bladeNum* is not operational, then, according to this algorithm,

beginFailedBlade is set as *bladeNum*  
 endFailedBlade is set as *bladeNum*  
 failedPortMask is set as 0xff  
 failedPortValue is set as *portNum*

[0032] The ports in DMU A have last (least significant) 2 bits set to 00. The ports in DMU B have last 2 bits set to 01. The ports in DMU C have last 2 bits set to 10 and the ports in DMU D have last 2 bits set to 11. In an example scenario when all the ports in DMU C fail in blade numbered *bladeNum*, then, according to this algorithm,

beginFailedBlade is set as *bladeNum*  
 endFailedBlade is set as *bladeNum*  
 failedPortValue is set as 0b 0000 0010 and  
 failedPortMask is set as 0b 0000 0011

[0033] While the invention has been particularly shown and described with respect to illustrative and preformed embodiments thereof, it will be understood by those skilled in the art that the foregoing and other changes in form and details may be made therein without departing from the spirit and scope of the invention which should be limited only by the scope of the appended claims.

### 30 Claims

1. In a networking environment comprising one or more network processing (NP) devices for routing data packets from a source to a destination via a switch fabric, with each network processing device supporting a number of interface ports, a system for ensuring packet routing from one network processing device to a target network processing device via a target interface port, said system comprising:

mechanism for tracking operational status of each network processor device and operational status of destination ports supported by each said network processor device in said system, said operational status being maintained at each network processing device;

40 said network processor devices including mechanism for determining the operational status of a target network processing device and target interface port of a current packet to be routed prior to said routing, routing mechanism for routing packets from source NP devices to destination NP devices and destination ports thereof in accordance with a packet routing protocol, said routing mechanism routing said current packet to a target network processor device and destination port when said target network processor device and destination ports thereof are determined as operational, and routing packets to another operational NP device and port thereof upon determination of non-operational target network processor device and destination port, whereby proper routing of packets is guaranteed with minimum packet lost.

2. The system for ensuring packet routing in accordance with Claim 1, wherein said routing mechanism implements an Equal Cost Multi-Path ECMP protocol including next hop routing table for mapping a destination address associated with a packet to be forwarded to one or more next hop options in said networking environment.

55 3. The system for ensuring packet routing in accordance with Claim 1, wherein each network processor device maintains a data structure receiving values from said tracking mechanism indicating status of said network processor devices, said determining mechanism implementing logic for comparing said received value against a first value indicating all NP devices are operational prior to routing of a current packet, and initiating routing of said packet to said target when said values match.

5	4. The system for ensuring packet routing in accordance with Claim 3, wherein said first value is a second value comparing a particular NP device that is not operational, said determining mechanism implements a bit representation of said packet routing logic for a target NP device that is not operational, said determining mechanism logic for said target NP device is not operational, and initializing routing logic of said packet to another NP device when said target NP device is not operational.
10	5. The system for ensuring packet routing in accordance with Claim 1, wherein each value includes a set of mask bits and a set of bits representing said target destination port, said determining mechanism logic for said target destination port is not operational, and initializing routing logic of said packet to said target destination port when said target destination port is not operational.
15	6. The system for ensuring packet routing in accordance with Claim 5, wherein said data structure receives two values resulting of said packet to said NP device and target port when said values match.
20	7. The system for ensuring packet routing in accordance with Claim 5, wherein said data structure receives two values defining a bit representation of a target NP device of a packet to be routed against said each of said two values defining said target destination port bits, and initializing re-routing of said packet to another destination port outside said range when said result does not match said target destination port bits.
25	8. A method for ensuring packet routing in a networking environment comprising one or more network processing (NP) devices for routing data packets from a source to a destination via a switch fabric, with each network processing device supporting said network processor device in said system, and maintaining said operational ports of each network processing device.
30	a) tracking operational status of each network processor device in said system, and maintaining said operational ports of each network processing device;
35	b) determining the operational status of a target network processing device and target interface port of a current packet to be routed prior to said routing at a current NP device; and,
40	c) routing packets from source NP devices to destination NP devices and destination ports thereof in accordance with a current routing protocol, a current packet being routed to a target network processor device and destination port when said target network processor device and destination ports thereof are forwarded to one or more next hop options in said networking environment.
45	d) determining port when said target network processor device and destination port thereof upon determining NP devices to destination NP devices and destination ports thereof is in accordance with Equal Cost Multi-Path (ECMP) protocol, said routing step (c) including mapping a destination address from said tracking step indicating final status including a data structure for receiving values determined from said tracking step indicating the step of ensuring packet routing in accordance with Claim 8, wherein said step of maintaining said opera-
50	10. The method for ensuring packet routing in accordance with Claim 8, wherein said step of maintaining said opera-
55	11. The method for ensuring packet routing in accordance with Claim 10, wherein said determined value is a second value comparing a particular NP device that is not operational, said determining mechanism logic for said target NP device is not operational, and initializing routing logic of said packet to another NP device when said target NP device is not operational.
60	12. The method for ensuring packet routing in accordance with Claim 11, wherein said received value is a second value comparing a particular NP device that is not operational, said determining mechanism logic for said target NP device is not operational, and initializing routing logic of said packet to another NP device when said target NP device is not operational.

13. The method for ensuring packet routing in accordance with Claim 10, wherein said determining step b) includes the step of implementing logic for comparing said received value against a first value indicating all interface ports for said NP devices are operational prior to routing of a current packet, and initiating routing of said packet to said NP device and target port when said values match.

5 14. The method for ensuring packet routing in accordance with Claim 13, wherein said first value includes a set of mask bits and a set of bits representing said target destination port, said determining step b) including the step of implementing bitwise logic for comparing said received value against said mask bit set and obtaining a first result, comparing said first result against said target destination port bits, and initiating re-routing of said packet to another destination port when said first result does not match said target destination port bits.

10 15. The method for ensuring packet routing in accordance with Claim 13, wherein said data structure receives two values defining a range of NP devices that are not operational, said determining step b) implementing logic for comparing a bit representation of a target NP device of a packet to be routed against said each of said two values defining said range, and initiating re-routing of said packet to another destination port outside said range when said bit representation of said target NP device falls within said two values.

15 16. A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform method steps for ensuring packet routing in a networking environment comprising one or more network processing (NP) devices for routing data packets from a source to a destination via a switch fabric, with each network processing device supporting a number of interface ports, said method steps comprising:

20 a) tracking operational status of each network processor device and operational status of destination ports supported by each said network processor device in said system, and maintaining said operational status at each network processing device;

25 b) determining the operational status of a target network processing device and target interface port of a current packet to be routed prior to said routing at a current NP device; and,

30 c) routing packets from source NP devices to destination NP devices and destination ports thereof in accordance with a packet routing protocol, a current packet being routed to a target network processor device and destination port when said target network processor device and destination ports thereof are determined as operational, or being routed to another operational NP device and port thereof upon determination of non-operational target network processor device and destination port, whereby proper routing of packets is guaranteed with minimum packet lost.

35 17. The program storage device readable by a machine in accordance with Claim 16, wherein said routing of packets from source NP devices to destination NP devices and destination ports thereof is in accordance with Equal Cost Multi-Path (ECMP) protocol, said routing step c) including mapping a destination address associated with a packet to be forwarded to one or more next hop options in said networking environment.

40 18. The program storage device readable by a machine in accordance with Claim 16, wherein said step of maintaining said operational status includes maintaining a data structure for receiving values determined from said tracking step indicating status of said network processor devices.

45 19. The program storage device readable by a machine in accordance with Claim 18, wherein said determining step b) includes the step of implementing logic for comparing a received value against a first value indicating all NP devices are operational prior to routing of a current packet, and initiating routing of said packet to said target when said values match.

50 20. The program storage device readable by a machine in accordance with Claim 19, wherein said received value is a second value representing a particular NP device that is not operational, said determining step b) including the step of implementing logic for comparing a bit representation of a target NP device of a packet to be routed against this received second value and initiating routing of said packet to another NP device when said target NP device is not operational.

55 21. The program storage device readable by a machine in accordance with Claim 20, wherein said determining step b) includes the step of implementing logic for comparing said received value against a first value indicating all interface ports for said NP devices are operational prior to routing of a current packet, and initiating routing of said packet to said NP device and target port when said values match.

23. The program storage device readable by a machine in accordance with Claim 18, wherein said data structure receives two values defining a range of NP devices that are not operational, said determining step b) implementing logic for comparing said bit representation of a target NP device of a packet to be routed against said each of said two values defining said range, and initiating re-routing of said packet to another destination port outside said range when said bit representation of said target NP device falls within said two values.

22. The program storage device readable by a machine in accordance with Claim 21, wherein said first value includes a set of mask bits and a set of bits representing said target destination port, said determining step b) including the step of implementing bitwise logic for comparing said received value against said mask bit set and obtaining a first result, comparing said first result against said target destination port bits, and initiating re-routing of said packet to another destination port when said first result does not match said target destination port bits.

FIG. 1




FIG. 2A




FIG. 2B

**50 ECMP Forwarding Table**

52	50	70	60a	60b	60c
Subnet	Action Data	Next Hop0	Next Hop1	Next Hop2	
7.***	(30%,80%)	9.1.1.1	8.1.1.1	6.1.1.1	
	72				




FIG. 3